Maritime Situational Awareness in the era of Large Language/Reasoning Models

Alexander Artikis[1,2]
Andreas Kouvaras[1]

[1]University of Piraeus, Greece
[2]NCSR Demokritos, Athens, Greece

https://cer.iit.demokritos.gr

Tutorial Resources

- Slides: https://cer.iit.demokritos.gr/blog/talks/maris25/
- Papers, code, datasets: http://cer.iit.demokritos.gr

Complex Event Recognition (CER)*,†

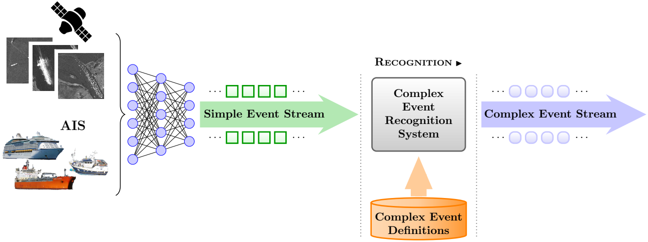

* Giatrakos et al, Complex Event Recognition in the Big Data Era: A Survey. VLDB Journal, 2020.
† Alevizos et al, Probabilistic Complex Event Recognition: A Survey. ACM Computing Surveys, 2017.

Complex Event Recognition (CER)*,†

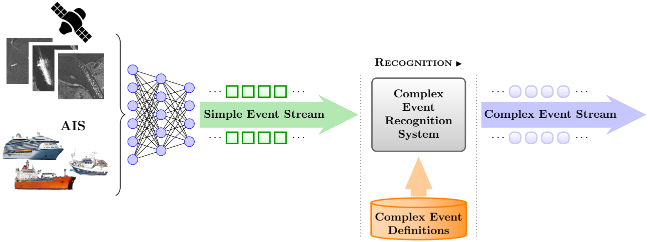

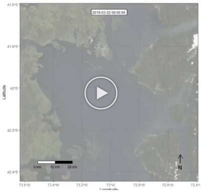

https://rdcu.be/cNkQE

*Giatrakos et al, Complex Event Recognition in the Big Data Era: A Survey. VLDB Journal, 2020.

†Alevizos et al, Probabilistic Complex Event Recognition: A Survey. ACM Computing Surveys, 2017.

Maritime Situational Awareness*

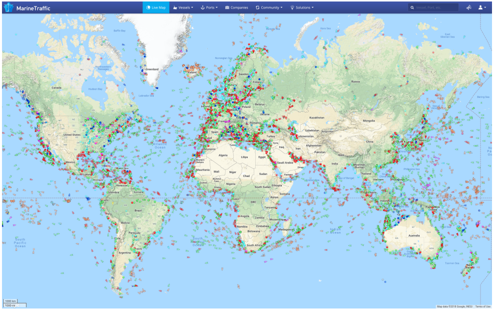

http://www.marinetraffic.com

*Artikis and Zissis, Guide to Maritime Informatics, Springer, 2021.

Maritime Situational Awareness*

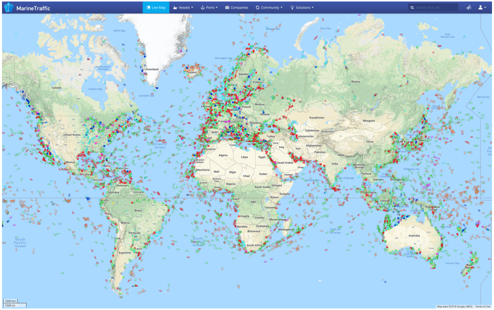

http://www.marinetraffic.com

https://cer.iit.demokritos.gr (fishing vessel)

*Artikis and Zissis, Guide to Maritime Informatics, Springer, 2021.

Maritime Situational Awareness*

https://cer.iit.demokritos.gr (tugging)

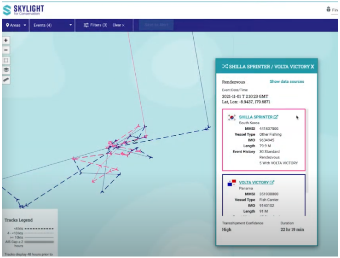

https://www.skylight.global (rendez-vous)

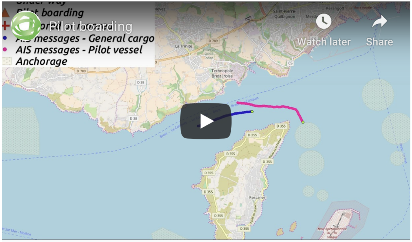

https://cer.iit.demokritos.gr (pilot boarding)

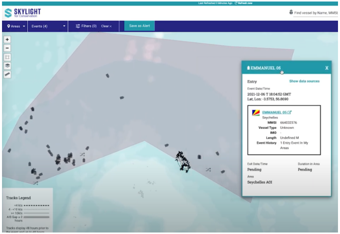

https://www.skylight.global (enter area)

*Artikis and Zissis, Guide to Maritime Informatics, Springer, 2021.

Data Challenges

- Velocity, Volume: Millions of position signals/day at European scale.

Data Challenges

- Velocity, Volume: Millions of position signals/day at European scale.
- Variety: Position signals need to be combined with other data streams
  - Weather forecasts, sea currents, etc.
- ... and static information
  - NATURA areas, shallow waters areas, coastlines, etc.

Data Challenges

- Velocity, Volume: Millions of position signals/day at European scale.
- Variety: Position signals need to be combined with other data streams
  - Weather forecasts, sea currents, etc.
- ... and static information
  - NATURA areas, shallow waters areas, coastlines, etc.
- Lack of Veracity: GPS manipulation, vessels reporting false identity, communication gaps.

Data Challenges

- Velocity, Volume: Millions of position signals/day at European scale.
- Variety: Position signals need to be combined with other data streams
  - Weather forecasts, sea currents, etc.
- ... and static information
  - NATURA areas, shallow waters areas, coastlines, etc.
- Lack of Veracity: GPS manipulation, vessels reporting false identity, communication gaps.
- Distribution: Vessels operating across the globe.

Requirements

- Expressive representation
  - to capture complex relationships between the events that stream into the system.

Requirements

- Expressive representation
  - to capture complex relationships between the events that stream into the system.
- Efficient reasoning
  - to support real-time decision-making in large-scale, (geographically) distributed applications.

Requirements

- Expressive representation
  - to capture complex relationships between the events that stream into the system.
- Efficient reasoning
  - to support real-time decision-making in large-scale, (geographically) distributed applications.
- Automated knowledge construction
  - to avoid the time-consuming, error-prone manual CE definition development.

Requirements

- Expressive representation
  - to capture complex relationships between the events that stream into the system.
- Efficient reasoning
  - to support real-time decision-making in large-scale, (geographically) distributed applications.
- Automated knowledge construction
  - to avoid the time-consuming, error-prone manual CE definition development.
- Reasoning under uncertainty
  - to deal with various types of noise.

Requirements

- Expressive representation
  - to capture complex relationships between the events that stream into the system.
- Efficient reasoning
  - to support real-time decision-making in large-scale, (geographically) distributed applications.
- Automated knowledge construction
  - to avoid the time-consuming, error-prone manual CE definition development.
- Reasoning under uncertainty
  - to deal with various types of noise.
- Complex event forecasting
  - to support proactive decision-making.

Course Structure

- Part I: Introduction to Complex Event Recognition (CER).

Course Structure

- Part I: Introduction to Complex Event Recognition (CER).
- Part II: A system meeting the CER requirements ...

Course Structure

- Part I: Introduction to Complex Event Recognition (CER).
- Part II: A system meeting the CER requirements ...
- ... powered by Large Language/Reasoning Models (LLMs/LRMs) — Part III.

Course Structure

- Part I: Introduction to Complex Event Recognition (CER).
- Part II: A system meeting the CER requirements ...
- ... powered by Large Language/Reasoning Models (LLMs/LRMs) — Part III.
- Part IV: Open issues & further research.

Course Structure

- Part I: Introduction to Complex Event Recognition (CER).
- Part II: A system meeting the CER requirements ...
- ... powered by Large Language/Reasoning Models (LLMs/LRMs) — Part III.
- Part IV: Open issues & further research.

- Each of the first 3 parts is followed by tutorial questions (15').
- The top-3 students will be announced at the end of the course!

Introduction to Complex Event Recognition (CER)

Complex Event Recognition vs DataBase Management Systems*

Complex event recognition (CER) systems:

- Process data without storing them.

*Gugola and Margara, Processing Flows of Information: From Data Stream to Complex Event Processing. ACM Computing Surveys, 2012.

Complex Event Recognition vs DataBase Management Systems*

Complex event recognition (CER) systems:

- Process data without storing them.
- Data are continuously updated.
  - Data stream into the system in high velocity.
  - Data streams are large (usually unbounded).

*Gugola and Margara, Processing Flows of Information: From Data Stream to Complex Event Processing. ACM Computing Surveys, 2012.

Complex Event Recognition vs DataBase Management Systems*

Complex event recognition (CER) systems:

- Process data without storing them.
- Data are continuously updated.
  - Data stream into the system in high velocity.
  - Data streams are large (usually unbounded).
- No assumption can be made on data arrival order.

* Gugola and Margara, Processing Flows of Information: From Data Stream to Complex Event Processing. ACM Computing Surveys, 2012.

Complex Event Recognition vs DataBase Management Systems*

Complex event recognition (CER) systems:

- Process data without storing them.
- Data are continuously updated.
  - Data stream into the system in high velocity.
  - Data streams are large (usually unbounded).
- No assumption can be made on data arrival order.
- Users install standing/continuous queries:
  - Queries deployed once and executed continuously until removed.
  - Online reasoning.

*Gugola and Margara, Processing Flows of Information: From Data Stream to Complex Event Processing. ACM Computing Surveys, 2012.

Complex Event Recognition vs DataBase Management Systems*

Complex event recognition (CER) systems:

- Process data without storing them.
- Data are continuously updated.
  - Data stream into the system in high velocity.
  - Data streams are large (usually unbounded).
- No assumption can be made on data arrival order.
- Users install standing/continuous queries:
  - Queries deployed once and executed continuously until removed.
  - Online reasoning.
- Latency requirements are very strict.

---

*Gugola and Margara, Processing Flows of Information: From Data Stream to Complex Event Processing. ACM Computing Surveys, 2012.

Complex Event Recognition vs Deep Learning

We have Deep Learning and it seems to work. Can we go home?

Complex Event Recognition vs Deep Learning

We have Deep Learning and it seems to work. Can we go home?

CER:

- Formal semantics* for trustworthy models.

---

*Grez et al, A Formal Framework for Complex Event Recognition. ACM Transactions on Database Systems, 2021.

Complex Event Recognition vs Deep Learning

We have Deep Learning and it seems to work. Can we go home?

CER:

- Formal semantics* for trustworthy models.
- Explanation — why did we detect a complex event?

*Grez et al, A Formal Framework for Complex Event Recognition. ACM Transactions on Database Systems, 2021.

Complex Event Recognition vs Deep Learning

We have Deep Learning and it seems to work. Can we go home?

CER:

- Formal semantics* for trustworthy models.
- Explanation — why did we detect a complex event?
- Machine Learning is necessary. But:
  - Complex events are rare.
  - Supervision is scarce.

*Grez et al, A Formal Framework for Complex Event Recognition. ACM Transactions on Database Systems, 2021.

Complex Event Recognition vs Deep Learning

We have Deep Learning and it seems to work. Can we go home?

CER:

- Formal semantics* for trustworthy models.
- Explanation — why did we detect a complex event?
- Machine Learning is necessary. But:
  - Complex events are rare.
  - Supervision is scarce.
- More often than not, background knowledge is available — let's use it!

*Grez et al, A Formal Framework for Complex Event Recognition. ACM Transactions on Database Systems, 2021.

Complex Event Recognition vs Large Language/Reasoning Models

- LLMs/LRMs cannot be used (yet) for CER[*,†].

[*] Ishay and Lee, LLM+AL: Bridging Large Language Models and Action Languages for Complex Reasoning About Actions. AAAI 2025.

[†] Shojaee et al, The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity. https://ml-site.cdn-apple.com/papers/the-illusion-of-thinking.pdf, 2025.

Complex Event Recognition vs Large Language/Reasoning Models

- LLMs/LRMs cannot be used (yet) for CER[*,†].
- But they may support it — see Part III of the course!

[*] Ishay and Lee, LLM+AL: Bridging Large Language Models and Action Languages for Complex Reasoning About Actions. AAAI 2025.

[†] Shojaee et al, The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity. https://ml-site.cdn-apple.com/papers/the-illusion-of-thinking.pdf, 2025.

A Simple Unifying Event Algebra

$$
\begin{array}{rlll}
ce ::= & se & | & \\
& ce_1 \; ; \; ce_2 & | & \textit{Sequence} \\
& ce_1 \; \vee \; ce_2 & | & \textit{Disjunction} \\
& ce^* & | & \textit{Iteration} \\
& \neg \; ce & | & \textit{Negation} \\
& \sigma_\theta(ce) & | & \textit{Selection} \\
& \pi_m(ce) & | & \textit{Projection} \\
& [ce]_{T_1}^{T_2} & & \textit{Windowing} \text{ (from } T_1 \text{ to } T_2\text{)}
\end{array}
$$

- *Sequence*: Two events following each other in time.

A Simple Unifying Event Algebra

$$
\begin{array}{rlll}
ce ::= & se & | & \\
& ce_1 \; ; \; ce_2 & | & \textit{Sequence} \\
& ce_1 \; \vee \; ce_2 & | & \textit{Disjunction} \\
& ce^* & | & \textit{Iteration} \\
& \neg \; ce & | & \textit{Negation} \\
& \sigma_\theta(ce) & | & \textit{Selection} \\
& \pi_m(ce) & | & \textit{Projection} \\
& [ce]_{T_1}^{T_2} & & \textit{Windowing} \text{ (from } T_1 \text{ to } T_2\text{)}
\end{array}
$$

- *Sequence*: Two events following each other in time.
- *Disjunction*: Either of two events occurring, regardless of temporal relations.

A Simple Unifying Event Algebra

$$
\begin{array}{rll}
ce ::= & se & | \\
& ce_1 \ ; \ ce_2 & | \ \textit{Sequence} \\
& ce_1 \ \vee \ ce_2 & | \ \textit{Disjunction} \\
& ce^* & | \ \textit{Iteration} \\
& \neg \ ce & | \ \textit{Negation} \\
& \sigma_\theta(ce) & | \ \textit{Selection} \\
& \pi_m(ce) & | \ \textit{Projection} \\
& [ce]_{T_1}^{T_2} & \ \ \textit{Windowing} \text{ (from } T_1 \text{ to } T_2\text{)}
\end{array}
$$

- *Sequence*: Two events following each other in time.
- *Disjunction*: Either of two events occurring, regardless of temporal relations.
- The combination of *Sequence* and *Disjunction* expresses Conjunction (both events occurring).

A Simple Unifying Event Algebra

$$
\begin{array}{rll}
ce ::= & se & | \\
& ce_1 \; ; \; ce_2 & | \; \textit{Sequence} \\
& ce_1 \; \vee \; ce_2 & | \; \textit{Disjunction} \\
& ce^* & | \; \textit{Iteration} \\
& \neg \; ce & | \; \textit{Negation} \\
& \sigma_\theta(ce) & | \; \textit{Selection} \\
& \pi_m(ce) & | \; \textit{Projection} \\
& [ce]_{T_1}^{T_2} & \; \textit{Windowing} \text{ (from } T_1 \text{ to } T_2\text{)}
\end{array}
$$

- *Iteration*: An event occurring N times in sequence, where $N \geq 0$. This operation is similar to the *Kleene star* operation in regular expressions, the difference being that *Kleene star* is unbounded.

A Simple Unifying Event Algebra

$$
\begin{array}{rll}
ce ::= & se & | \\
& ce_1 \ ; \ ce_2 & | \ \textit{Sequence} \\
& ce_1 \ \vee \ ce_2 & | \ \textit{Disjunction} \\
& ce^* & | \ \textit{Iteration} \\
& \neg \ ce & | \ \textit{Negation} \\
& \sigma_\theta(ce) & | \ \textit{Selection} \\
& \pi_m(ce) & | \ \textit{Projection} \\
& [ce]_{T_1}^{T_2} & \ \ \textit{Windowing} \text{ (from } T_1 \text{ to } T_2\text{)}
\end{array}
$$

- *Negation*: Absence of event occurrence.

A Simple Unifying Event Algebra

$$
\begin{array}{rlll}
ce ::= & se & | & \\
& ce_1 \ ; \ ce_2 & | & \textit{Sequence} \\
& ce_1 \ \vee \ ce_2 & | & \textit{Disjunction} \\
& ce^* & | & \textit{Iteration} \\
& \neg \ ce & | & \textit{Negation} \\
& \sigma_\theta(ce) & | & \textit{Selection} \\
& \pi_m(ce) & | & \textit{Projection} \\
& [ce]_{T_1}^{T_2} & & \textit{Windowing} \text{ (from } T_1 \text{ to } T_2)
\end{array}
$$

- *Negation*: Absence of event occurrence.
- *Selection*: Select those events whose attributes satisfy a set of predicates/relations θ, temporal or otherwise.

A Simple Unifying Event Algebra

$$
\begin{array}{rll}
ce ::= & se & | \\
& ce_1 \ ; \ ce_2 & | \ \textit{Sequence} \\
& ce_1 \ \vee \ ce_2 & | \ \textit{Disjunction} \\
& ce^* & | \ \textit{Iteration} \\
& \neg \ ce & | \ \textit{Negation} \\
& \sigma_\theta(ce) & | \ \textit{Selection} \\
& \pi_m(ce) & | \ \textit{Projection} \\
& [ce]_{T_1}^{T_2} & \ \ \textit{Windowing} \text{ (from } T_1 \text{ to } T_2\text{)}
\end{array}
$$

- *Projection*: Return an event whose attribute values are a possibly transformed subset of the attribute values of its sub-events.

A Simple Unifying Event Algebra

$$
\begin{array}{rll}
ce ::= & se & | \\
& ce_1 \; ; \; ce_2 & | \; \textit{Sequence} \\
& ce_1 \; \vee \; ce_2 & | \; \textit{Disjunction} \\
& ce^* & | \; \textit{Iteration} \\
& \neg \; ce & | \; \textit{Negation} \\
& \sigma_\theta(ce) & | \; \textit{Selection} \\
& \pi_m(ce) & | \; \textit{Projection} \\
& [ce]_{T_1}^{T_2} & \quad \textit{Windowing} \text{ (from } T_1 \text{ to } T_2\text{)}
\end{array}
$$

- *Projection*: Return an event whose attribute values are a possibly transformed subset of the attribute values of its sub-events.
- *Windowing*: Evaluate the conditions of an event pattern within a specified time window.

Processing Model

Selection strategies filter the set of matched patterns.

- Assume the pattern $\alpha ; \beta$ and the stream $(\alpha, 1),\ (\alpha, 2),\ (\beta, 3)$.

Processing Model

Selection strategies filter the set of matched patterns.

- Assume the pattern $\alpha;\beta$ and the stream $(\alpha, 1),\ (\alpha, 2),\ (\beta, 3)$.
- The multiple selection strategy produces $(\alpha, 1),\ (\beta, 3)$ and $(\alpha, 2),\ (\beta, 3)$.

Processing Model

Selection strategies filter the set of matched patterns.

- Assume the pattern $\alpha;\beta$ and the stream $(\alpha,1),\ (\alpha,2),\ (\beta,3)$.
- The multiple selection strategy produces $(\alpha,1),\ (\beta,3)$ and $(\alpha,2),\ (\beta,3)$.
- The single selection strategy produces either $(\alpha,1),\ (\beta,3)$ or $(\alpha,2),\ (\beta,3)$.
- The single selection strategy represents a family of strategies, depending on the matches actually chosen among all possible ones.

Instantaneous vs Interval-based Reasoning[*,†]

Consider:

- the pattern $\beta; (\alpha; \gamma)$
- and the stream $(\alpha, 1), (\beta, 2), (\gamma, 3)$.

Does the stream match the pattern?

[*] Paschke, ECA-LP / ECA-RuleML: A Homogeneous Event-Condition-Action Logic Programming Language. RuleML, 2006.

[†] White et al, What is "Next" in Event Processing?, PODS, 2007.

SASE*: Example (1)

PATTERN SEQ(*gapStart a*, *gapEnd b*, *speedChange c*)
WHERE partition-contiguity
AND *vesselId*
AND *c.velocity* $>$ *20*
WITHIN 3600

*Zhang et al. On complexity and optimization of expensive queries in complex event processing. SIGMOD 2014.

SASE*: Example (1)

PATTERN SEQ(*gapStart a*, *gapEnd b*, *speedChange c*)
WHERE partition-contiguity
AND *vesselId*
AND *c.velocity* $>$ *20*
WITHIN 3600

Quickly moving away from an area of suspicious activity:

- After a communication gap, ...
- a vessel changes speed to over 20 knots.
- Partition contiguity ensures that a, b, c refer to the same vessel (*vesselId*) and are contiguous with respect to that vessel.

---

* Zhang et al. On complexity and optimization of expensive queries in complex event processing. SIGMOD 2014.

SASE*: Example (2)

PATTERN SEQ($lowSpeedStart\ a,\ turn+\ b,\ lowSpeedEnd\ c$)
WHERE skip-till-next-match
AND *vesselId*
AND $b[i].heading - b[i-1].heading > 90$
WITHIN 21600

---

*Zhang et al. On complexity and optimization of expensive queries in complex event processing. SIGMOD 2014.

SASE*: Example (2)

PATTERN SEQ($lowSpeedStart\ a,\ turn+\ b,\ lowSpeedEnd\ c$)
WHERE skip-till-next-match
AND *vesselId*
AND $b[i].heading - b[i-1].heading > 90$
WITHIN 21600

Fishing pattern:

- A vessel slows down, ...
- begins a series of turns, where, for each pair of successive turns, their difference in heading is more than 90 degrees, ...
- and subsequently the vessel stops moving at a low speed.

*Zhang et al. On complexity and optimization of expensive queries in complex event processing. SIGMOD 2014.

Summary

What we've seen so far:

- Complex Event Recognition (CER) for Maritime Situational Awareness.
  - Related research.
  - Event algebras for CER.

Summary

What we've seen so far:

- Complex Event Recognition (CER) for Maritime Situational Awareness.
  - Related research.
  - Event algebras for CER.
- Requirements:
  - Expressive representation.
  - Efficient reasoning.
  - Automated knowledge construction.
  - Reasoning under uncertainty.
  - Complex event forecasting.

Summary

What we've seen so far:

- Complex Event Recognition (CER) for Maritime Situational Awareness.
    - Related research.
    - Event algebras for CER.
- Requirements:
    - Expressive representation.
    - Efficient reasoning.
    - Automated knowledge construction.
    - Reasoning under uncertainty.
    - Complex event forecasting.

Next: A system addressing the CER requirements.